The main point of this section is that there is an **excess burden**, sometimes referred to as a **deadweight loss**, to taxation. The main point is **not** that taxes will hurt the people who bear the taxes. That point is true, but obvious. What may be less obvious — but is true — is that the magnitude of this damage, measured in dollars, is greater than the benefit.

What's the benefit of taxation? The government raises some revenue from the tax. The excess burden is the difference between the cost in dollars of this tax to the taxpayers, and the magnitude of this revenue yield. Every dollar raised by the government turns out to cost more than one dollar to the people who bear the tax.

To put the point another way, suppose that the government, for some reason, chose to impose a tax on some good, and chose as well to **compensate** people for the damage done to them by the tax. How much would they have to compensate them for this damage? If the tax increases the total price that they pay for some commodity, then that is exactly what the compensating variation measures. In this section I attempt to show that this compensating variation will exceed the tax revenue actually raised by the tax.

Yet another way to put the same basic point : The tax harms the people who bear it ; the tax puts them on lower indifference curves. Suppose instead of putting in the tax, the government simply took some money from people, exactly the amount of money which would harm them as much as the tax. By definition, that's the **equivalent variation** to putting in the tax. Another way of phrasing the basic point of this section is that the government would be better off taking the money in this fashion : the equivalent variation to the tax is bigger than the revenue raised by the tax.

Ignoring everything done in the previous section, on tax incidence, suppose that an excise tax on some commodity is born entirely by buyers of the commodity. That is, assume now that the government puts in a unit tax of t on the commodity, and the total price $P = p + t$ paid by buyers goes up by exactly t. Then

Result : The compensating variation to the price increase caused by the tax is **greater** than the revenue raised by the tax.

The proof of this result follows from Figure 3. Here the taxed good is Y, the consumption of which is measured on the vertical axis. To make matters more simple, I have chosen units so that the price of good X is \$1 per unit. (That is, good X is the **numéraire**.) That way, movements along the **horizontal** axis in Figure 3 are measured in dollars. There are two numbers which must be calculated. Both are measured in dollars. One number is the tax revenue collected by the government, and the other is the amount of compensation it would have to pay someone for the damage done by the tax.

The second number can be illustrated, in Figure 3, in the following manner. The tax–induced increase in the price of good Y pivots the budget line, making it less steep. This moves the person to a lower indifference curve. Compensating the person shifts out the new, less steep budget line, until the consumer is just as well off as before. She is just as well off as before when her indifference curve which is tangent to the new, shifted–out, less steep budget line is the same as the indifference curve tangent to the original budget line. The magnitude of the compensation is the horizontal shift in the budget line.

So the original budget line for the consumer is the inner of the two steep lines in figure 1, the budget line that goes through the consumption bundle A. (It's the green dashed line labeled "budget line : no tax" in the figure.) A is the consumption bundle which the consumer would choose, in the original situation, in which there is no tax on good Y (and in which she receives no compensation). The tax on good Y would pivot the budget line in, around the point $(80, 0)$, and make the consumer worse off. (That is the dark blue line labeled "budget line : tax on clothing" in the figure.) Parallel to this is the budget line the consumer would face if good Y were taxed, and as well she were compensated for the damage. This new budget line is the less steep line through "B" in figure 1 (the [purple] one labeled "budget line : after compensation"). The tax on good Y makes the budget line less steep, since the slope of a budget line is $-P_X/P_Y$. The compensation shifts out the budget line, until it is just tangent to the indifference curve through A, representing the level of utility the consumer had originally, before any taxes or compensation. If she did face this price increase for good Y, and did get compensated, then she would choose consumption bundle B, where the indifference curve is tangent to the new budget line.

Notice that the consumption point that she would choose if good Y is taxed, and if she were compensated, must be **below and to the right** of her original consumption point A in the diagram. That's because she's on the same indifference curve, but her budget line is less steep. Given the shape of the budget line, and the fact that the indifference curve displays a declining marginal rate of substitution, she must have to move down and the right to be tangent to the new, less steep, budget line.

How much compensation would she have to receive, after the tax, to get back to her original indifference curve? In the figure, that is measured by the amount that the less steep after–tax budget line was shifted right (the distance between the dark blue "budget line : tax on clothing" and the purple "budget line : after compensation"). Remember that the price of good X, the numéraire, is \$1. That means that the horizontal intercept of a budget line is exactly the income she has : it's the amount of X she could buy if she spent all her income on X. (If the price of good X were not \$1, all I would have to do is divide everything by the price of good X.)

In Figure 3, that horizontal distance is actually 40. (That is, the dark blue line has a horizontal intercept of 80, and the purple one an intercept of 120.)

Next, how much tax revenue does the government raise? To read this from the diagram, the following result is needed : lines parallel to the original (no tax) budget line measure the **cost** of consumption bundles, exclusive of tax. That is, what is the cost of a consumption bundle $(\bar{x}, \bar{y})$, using the net–of–tax prices? It's $C = p_x\bar{x} + p_y\bar{y}$, where p_x and p_y are the net–of–tax prices of the commodities. So any bundle (x, y) such that

$$p_x x + p_y y = p_x \bar{x} + p_y \bar{y} \tag{cost}$$

has exactly the same net–of–tax cost as the bundle $(\bar{x}, \bar{y})$. Equation (*cost*) defines a line, through the consumption bundle $(\bar{x}, \bar{y})$, with slope $-p_x/p_y$: a line through a consumption bundle, with slope equal to the slope of the original "no tax" budget line. The actual value of the net–of–tax cost of this bundle can be read off the diagram, from where this line hits the horizontal axis (because the price of good X, p_x, equals \$1).

In Figure 3, the net–of–tax cost of the bundle B can be measured in this way, by drawing a line through B with the same slope as the original "no tax" budget line. (That is the light blue dot–dashed line labeled "cost of bundle B" in the figure.) In Figure 3, that cost equals 86.67.

What tax revenue does the government collect from this person? If good Y is taxed, and if she is compensated for the damage, then she chooses the consumption bundle $B \equiv (x_B, y_B)$. (In the figure $(x_B, y_B) = (60, 60)$.) The net–of–tax cost of the bundle is $p_x x_B + p_y y_B$ (in the figure, \$86.67). What she actually pays for the bundle B is the cost of the bundle **including all taxes**, $p_x x_B + (p_y + t) y_B$, since here it is good Y which is being taxed. That tax–included cost is her total expenditure on all goods and service, that is, the income — including any compensation she has received — which she has to spend. It is the horizontal intercept of her budget line through B — after taxes, and after compensation. (In the figure, this income is the horizontal intercept of the purple line, which is \$120.)

The tax revenue collected is the unit tax times the number of units purchased : $t y_B$. That is the difference between the tax–included cost $p_x x_B + (p_y + t) y_B$, and the net–of–tax cost $p_x x_B + p_y y_B$.

It's the horizontal distance between the horizontal intercepts of two lines through B : the tax–included budget line with slope $-p_x/(p_y + t)$ and the net–of–tax cost line with slope $-p_x/p_y$.

Summarizing the graphical illustration :

i The compensation which must be paid is the horizontal distance between the budget line through B with slope reflecting tax–included prices $p_x/(p_y + t)$ and the consumer's original income. That's the same thing as the horizontal distance between the budget line through B with slope reflecting tax–included prices $p_x/(p_y + t)$ and the budget line through A with slope p_x/p_y.

ii The tax collected is the horizontal distance between the budget line through B with slope reflecting tax–included prices $p_x/(p_y + t)$ and the budget line through B with slope reflecting net–of–tax prices p_x/p_y.

In Figure 3 the compensation is the horizontal distance between 120 and 80, or \$40. The tax revenue collected is the horizontal distance between 120 and 86.67, or 33.33. The compensation paid exceeds the revenue collected. There is an **excess burden**, or **deadweight loss** of $40 - 33.33 = 6.67$ from the tax on good Y.

Comparing measures i and ii above, note that both involve subtracting something from the consumer's income after compensation (the intercept of the budget line through B with slope $-p_x/(p_y+t)$). The compensation subtracts off the net–of–tax cost of bundle A ; the tax revenue subtracts off the net–of–tax cost of bundle B. In the example in the figure, the compensation is larger than the tax revenue because the net–of–tax cost of bundle B is larger than the net–of–tax cost of bundle A.

That is, the line (with slope $-p_x/p_y$) through B is to the right of the line (with slope $-p_x/p_y$) through A. But that is the same thing as saying that B is to the right of A. In other words

whenever the original consumption bundle (when good Y is not taxed) A is to the left of the "new" consumption bundle (when good Y is taxed, and when the consumer is compensated) B, then the compensation paid will exceed the tax revenue collected— by the horizontal distance between the bundles

But the bundle B **must** be to the right of the bundle A, whenever a tax is put on good Y. Bundles A and B are both on the same indifference curve. Bundle A is the tangency of that indifference curve with a line of slope $-p_x/p_y$. Bundle B is the tangency with a line of slope $-p_x/(p_y+t)$. As long as $t>0$, then that second line is less steep than the first. As long as the person's indifference curves have the usual shape, getting less steep as we move down and to the right, then the tangency with a less steep line must occur to the right of the tangency with a more steep line. Therefore

General result : The cost to the consumer of a tax (measured by the income required to compensate for its damage) must always exceed the revenue raised by the tax.

The only way that the compensating to the tax will not exceed the revenue from the tax will be if the points A and B in the diagram coincide. Since taxing the good Y must raise its price, the budget line after the tax and compensation must be less steep than the one before. Since indifference curves must get less steep as we move down them and to the right, then A cannot be to the right of B. Only if goods X and Y were **perfect complements**, so that the indifference curves were L–shaped, would A and B coincide. And only then would the excess burden be zero.

Rosen, Wen, and Snoddon analyze the excess burden using the **equivalent variation**, rather than the compensating variation. The question then is whether the government would be better off levying a tax on some good, or instead collecting from the consumer a lump sum amount equal to the equivalent variation to the tax. Notice, by definition, the two proposals are equally bad for the consumer : the equivalent variation is defined as the amount of money that the government can take away from the person which would leave the person just as badly off as having the tax imposed. The excess burden using this definition is defined as the difference between this equivalent variation, and the revenue raised by the tax. This difference is always positive, unless the two goods are perfect complements, in which case it is zero. Any substitution possibility by the consumer means that the excess burden is positive.

Figure 4 measures the excess burden of a tax, using the equivalent variation. The equivalent variation to a tax is the bribe that a consumer would be willing to pay in order to avoid the tax. In the figure, that bribe is the horizontal distance between her original budget line, and a line tangent to the indifference curve which she would wind up on if there was a tax. (It's the distance between the dashed green line "budget line : no tax", and the purple line "budget line : after bribe".) The tax revenue is the difference between the tax–included cost of the bundle she chooses (if there was a tax) and the net–of–tax cost of that bundle. In Figure 4, E is the bundle which she would choose if there were a tax. So the revenue raised from the tax is the distance between the horizontal intercepts of the two lines through E : the tax–included cost line with slope $-p_x/(p_y + t)$ (the dark blue line "budget line : tax on clothing") and the net–of–tax cost line with slope $-p_x/p_y$ (the dot–dashed light blue line "cost of bundle E").

Using the equivalent variation, the excess burden of the tax is the difference between these two amounts, which is the difference between the net–of–tax cost of her after–tax consumption bundle E, and the net–of–tax cost of the bundle D she would choose if she paid a bribe instead of the tax. Since her budget line is less steep when there is a tax, E must be to the right of D, so that this excess burden must be positive. But its value is not the same as the excess burden calculated using the compensating variation. (In figure 2, the excess burden, calculated using the equivalent variation to the tax increase, is \$4.44, smaller than the excess burden calculated in figure 1, using the compensating variation, of \$6.67.)

Although these measures of excess burden of a tax differ slightly in value, both are positive, and both get larger, the larger is the elasticity of substitution along the consumer's indifference curve.

One of the advantages of using the equivalent variation is that it suggests an alternative to distortionary commodity taxation : a lump–sum tax. A lump–sum tax is defined as any tax in which the amount collected does not depend on the person's behaviour. Instead of collecting a tax in which the amount collected depends on a person's clothing consumption, it is better to collect the revenue in a lump–sum amount. The material on the Edgeworth Box, and on efficiency, suggested that lump–sum taxation was efficient whereas commodity taxation was not. The definition of excess burden now gives a dollar measure of how much this inefficiency costs.

So the main point so far is that excise taxes have an excess burden, and that lump–sum taxation could raise more revenue at less damage to the taxpayers.

But why is lump–sum taxation not more widely used? This is a question addressed very briefly in the introductory section. The short answer is that we are not able to implement equitable lump–sum taxation. It is not the case that lump–sum taxation is impossible. A head tax (sometimes called a poll tax), in which every person pays the same tax is lump–sum, since what you pay does not depend in any way on your choices. But such a tax is very regressive : rich people and poor people pay the same tax. More relevant to the particular issue of replacing excise taxes, the result was that we could improve on an excise tax by charging **each person** a lump–sum tax equal to the equivalent variation of the cost **to her** of the excise tax. Notice that the concepts of equivalent and compensating variations are *individual* concepts. We would have to calculate a CV or EV for each person, depending on her own indifference curves. So the lump–sum tax which would constitute a Pareto improvement on the excise tax involves a separate lump–sum tax for each person, equal to her EV to the excise tax. People with bigger EV's would pay bigger taxes. So we would have to know what each person's tastes were in order to implement the lump–sum tax. Moreover, we cannot use information on people's actual consumption behaviour to help ascertain these tastes. Suppose I knew that I would face a lump–sum tax equal to my equivalent variation to a tax on gasoline. Suppose further that I knew that the government would try and calculate this EV using information on my preferences, obtained by monitoring my consumption behaviour. Then I would have an incentive to modify my behaviour, to make it appear that my EV was lower than it was. In other words, here the tax would not really be lump–sum, because the amount of the tax would depend on my behaviour. Only if we could base the tax on information obtained beforehand somehow, would the tax be lump–sum.

Using the Expenditure Function

The main result of this section, that the deadweight loss, or excess burden, is positive, follows much more easily using the expenditure function.

What is the compensating variation to the price increase caused by the tax on good X? It is the increase in the cost of the consumer getting to her pre–tax level of utility u_0.

$$CV = E(p_x + t_x, p_y, u_0) - E(p_x, p_y, u_0) \tag{17}$$

How much tax revenue would the government collect from a tax on good X? The tax per unit is t_x, so that the total tax revenue is the tax per kilo of food, times the number of kilos of food which the person chooses to consume. That's her demand for food.

$$TR = t_x X^H(p_x + t_x, p_y, u_0) \tag{18}$$

The excess burden EB is the difference between these numbers, $CV - TR$, or

$$EB = E(P_X, p_y, u_0) - E(p_x, p_y, u_0) - t_x X^H(P_X, p_y, u_0) \tag{19}$$

where

$$P_X \equiv p_x + t_x$$

is the price of food with the tax included. Naturally, there is no excess burden if there is no tax : expression (19) is zero when $t_x = 0$.

How does the excess burden change when the tax increases? Differentiating expression (19) with respect to the tax rate t_X,

$$\frac{\partial EB}{\partial t_x} = \frac{\partial E}{\partial P_X} - X^H(P_X, p_y, u_0) - t_x \frac{\partial X^H}{\partial P_X} \tag{20}$$

Shepard's Lemma says that $\frac{\partial E}{\partial P_X} = X^H(P_X, p_y, u_0)$: the first two terms in expression (20) must cancel each other out. So, from Shepard's Lemma, the change in the excess burden when the tax rate is increased on good X is

$$\frac{\partial EB}{\partial t_x} = -t_x \frac{\partial X^H(P_X, p_y, u_0)}{\partial P_X} \tag{21}$$

The last result in the previous section was that compensated demand curves must slope down : $\frac{\partial X^H}{\partial P_X} \leq 0$. So this results says that the derivative in equation (21) is non–negative — and only equals zero if there are kinks in the consumer's indifference curve. That is, the excess burden of a tax on good X starts at 0, when the tax is zero, and must increase with the tax rate as it becomes positive : the excess burden cannot be negative, and must be positive if there are no kinks in the indifference curves.

[*Notice that the tax revenue here is the unit tax rate* t_x, *times the quantity demanded,* $X^H(P_X, p_y, u_0)$. *That is the demand for the good if the tax of* t_x *is imposed, and if the person is compensated so as to keep her on the same indifference curve. That is, the excess burden here is the amount of money the government would lose, if it taxed a good, and then compensated the consumer for the damage inflicted by the tax.*]

If the equivalent variation is used instead of the compensating variation, then the excess burden still must be non–negative, although it will differ in value from the excess burden calculated using the compensating variation. If we simply took away enough money from the consumer, so that the damage done to her was the same as that done by the tax on good X, the amount of money we could take away (and still leave the consumer no worse off than she would be with the tax on good X) is

$$EV = E(P_X, p_y, u_1) - E(p_x, p_y, u_1) \tag{22}$$

where u_1 is the level of utility the consumer gets if good X is taxed, and if she is not compensated. The tax revenue collected (if the consumer is not compensated) is

$$TR = t_x X^H(P_X, p_y, u_1) \tag{23}$$

The excess burden is how much **more** money the government could collect if it simply took some money away from the consumer, rather than taxing her consumption of X — provided the two policies leave the consumer equally badly off. So, using the equivalent variation,

$$EV = E(P_X, p_y, u_1) - E(p_x, p_y, u_1) - t_x X^H(P_X, p_y, u_1) \tag{24}$$

Using Shephard's Lemma again, the change in the excess burden, measured this way, when the tax increases, is

$$\frac{\partial EB}{\partial t_x} = -t_x \frac{\partial X^H(P_X, p_y, u_1)}{\partial P_X} \geq 0 \tag{25}$$

(equalling zero only if there are kinks in the consumer's indifference curves). So the excess burden measured using the equivalent variation must start out at zero when the tax is 0, and must increase above 0 as the tax rate increases.

The excess burden of a tax of t_x on good X equals [exactly] the area ABC in Figure 5. That is the approximately–triangular area under the compensated demand curve (the red curve in Figure 5), between the heights of the price p_x without tax (the blue line) and the price $P_X = p_x + t_x$ including the tax (the green line).

The reason for this result : the cost of the tax (which is the CV, if the Hicksian demand curve corresponds to the pre–tax initial utility level u_0) is the area between the two lines, between the vertical axis and the compensated demand curve. That is, the CV is the area $ECBAD$ in Figure 5.

Equation (17) above defined the CV :

$$CV = E(p_x + t_x, p_y, u_0) - E(p_x, p_y, u_0) \tag{17}$$

Now recall that the partial derivative of the expenditure function is the Hicksian demand function : that was Shephard's lemma.

Now I use the fundamental theorem of calculus, which says that the integral is the reverse of the derivative. In particular, that

$$\int_a^b f'(z)dz = f(b) - f(a) \tag{26}$$

for any function $F(\cdot)$. That means that

$$\int_a^b \frac{\partial E(P_x, p_y, u_0)}{\partial P_X} dP_X = E(a, p_y, u_0) - E(b, p_y, u_0) \tag{27}$$

so that CV defined in equation (17) is the integral of the the derivative of $E(P_X, p_y, u_0)$ with respect to P_X, as P_X runs from p_x to $p_x + t_x$. So Shepard's Lemma then implies that CV is the integral of $X^H(P_X, p_y, u_0)$ as P_X runs from p_x to $p_x + t_x$.

So what I have so far : the cost of a tax increase is (exactly) equal to the integral of the consumer's Hicksian (compensated) demand function for the taxed good, between the before–tax and after–tax prices.

Now the integral of a function is the area under the graph of the function. But in economics, we put the price, the independent variable, on the vertical axis, and quantity (the dependent variable) on the horizontal axis. So the integral of the Hicksian demand function is the area "inside the demand function, between the before–tax price and the after–tax price. That's the area $ECBAD$ in Figure 5.

The CV is not the excess burden. It's the total burden. The excess burden (also known as the deadweight loss) is the total burden, minus the tax revenue, $CV - TR$. In Figure 5, the tax revenue, which is the tax rate per kilo, times the number of kilos demanded, equals the area of rectangle $ECDA$. The height of the rectangle is the tax t_x, and the width is the quantity demanded $X^H(p_x + t_x, p_y, u_0)$ when the good is taxed.

Subtracting the tax revenue — the rectangle $ECDA$, from the total CV — the area $ECBAD$, leaves the approximately–triangular area ABC as the exact measure of the deadweight loss of the tax.

Note that this measure is "approximately" triangular, because the compensated demand curved is not usually exactly a straight line. In Figure 5, the area ABC would be 3.75 if the (red) demand curve really were a straight line. It's not a straight line, and the actual exact value of the deadweight loss in Figure 5 is 3.5.